

Formal Verification of Scalable NonZero Indicators

Shao Jie Zhang[†], Yang Liu[†], Jun Sun[†], Jin Song Dong[†] and Yanhong A. Liu[‡]

[†]National University of Singapore

shaojiezhang@nus.edu.sg, {liuyang,sunj,dongjs}@comp.nus.edu.sg

[‡]State University of New York at Stony Brook

liu@cs.sunysb.edu

Abstract

Concurrent algorithms are notoriously difficult to design correctly, and high performance algorithms that make little or no use of locks even more so. In this paper, we describe a formal verification of a recent concurrent data structure Scalable NonZero Indicators. The algorithm supports incrementing, decrementing, and querying the shared counter in an efficient and linearizable way without blocking. The algorithm is highly non-trivial and it is challenging to prove the correctness. We have proved that the algorithm satisfies linearizability, by showing a trace refinement relation from the concrete implementation to its abstract specification. These models are specified in CSP and verified automatically using the model checking toolkit PAT.

1 Introduction

Concurrent algorithms are notoriously difficult to design correctly, and high performance algorithms that make little or no use of locks even more so. The main correctness criterion of the concurrent algorithm design is linearizability [7]. Informally, a shared object is *linearizable* if each operation on the object can be understood as occurring instantaneously at some point, called *linearization point*, between its invocation and its response, and its behavior at that point is consistent with the specification for the corresponding sequential execution of the operation.

Formal verification of linearizability is challenging because the correctness often relies on the knowledge of linearization points, which is difficult or even impossible to identify. These proofs are too long and complicated to do (and check) reliably “by hand”. Hence, it is important to develop techniques for mechanically performing, or at least checking, such proofs.

In this paper, we present an approach to verify linearizability based on refinement relations between abstract speci-

fication and concrete implementation models of a concurrent algorithm. Both are specified using an event-based modeling language, which has formal semantics based on labeled transition systems. We have used this approach to formally verify a recent concurrent algorithm Scalable NonZero Indicators (SNZI) due to Ellen et al. [6], since the algorithm as a complex and useful implementation serves a good candidate for automatic verification. Our approach also builds on earlier work [9] in which we proved (and in some cases disproved and/or improved) a number of non-blocking implementations of concurrent stacks, queues and dequeues. We have made considerable progress in understanding how to model algorithms including specifications and implementations to allow model checking to scale up and handle bigger cases. The complete model of SNZI algorithm is built inside a novel model checking tool, PAT [12] (<http://pat.comp.nus.edu.sg>).

The rest of the paper is structured as follows. Section 2 briefly introduces the SNZI algorithm. Section 3 gives the standard definition of linearizability. Section 4 shows how to express linearizability using refinement relations in general. Section 5 gives the SNZI model in our modeling language. Section 6 presents the verification and experimental results. Section 7 discusses related work and concludes.

2 The SNZI Algorithms

A SNZI object behaves similarly to traditional shared counter. It has one shared integer variable *surplus* and supports three operations: *Arrive* increments *surplus* by 1 when a process enters; *Depart* decrements *surplus* by 1 when the process leaves; the only difference from traditional counters is *Query* operation: it returns a boolean value indicating whether the value of *surplus* is greater than 0. We assume that each *Arrive* operation is always followed by a *Depart* operation for the same process. Therefore *surplus* is always greater or equal to 0. The pseudo code in Fig. 1 gives the specification of a SNZI object.

```

shared variable : Surplus : integer ; initially 0
bool Query() : return (Surplus > 0)
void Arrive() : Surplus ← Surplus + 1
void Depart() : Surplus ← Surplus - 1

```

Figure 1. SNZI Specification

In [6], the authors propose a rooted tree as the underlying data structure of the SNZI objects implementation. An operation on a child node may invoke operations on its parent. An important invariant is used to guarantee the correctness: the *surplus* of parent node is non-zero if and only if there exists at least one child whose *surplus* is non-zero. Thus, if the *surplus* of one node in the tree is non-zero, so does the root. A process begins *Arrive* operation on any node as long as the corresponding *Depart* will be invoked at the same node, and *Query* operation is directly invoked on the root. Every tree node has a counter X that is increased by *Arrive* and decreased by *Depart*. Since the operations on hierarchical nodes differ from those on root node, the algorithms are separated for hierarchical nodes and root node.

The code for hierarchical SNZI nodes is shown in Fig 2. An *Arrive* operation on a hierarchical node invokes *Arrive* operation on its parent node when increasing X from 0 to 1. Otherwise, it completes without invoking any operation. Moreover, a process which increases X from 0 to 1 should firstly set X by an intermediate value $\frac{1}{2}$. Any process which sees $\frac{1}{2}$ must help that process to invoke *parent.Arrive* and try to change X to 1. If a process succeeds in invoking *parent.Arrive* but fails in setting X to 1, it will invoke a compensating *parent.Depart*.

Similarly, a *Depart* operation on a hierarchical node only invokes *Depart* on its parent node when decreasing X from 1 to 0. A version number is added to X to ensure that every change of X will be detected in both *Arrive* and *Depart* operations for hierarchical nodes as well as root node.

The code for root node is shown in Fig 3. In order to reduce frequent accesses to X by *Query*, the solution for the root node separates out an indicator bit I from X . Hence every process can finish *Query* only by reading the bit I . The authors model all accesses to I using *Read*, *Write*, *Load Linked* and *Store Conditional* primitives to tolerate spurious failures when external applications try to modify I .

I is set to true after a 0 to 1 transition of X , and it is unset to false after a 1 to 0 transition of X . Furthermore, an announce bit a is added to X to indicate that I needs to be set. Similar to the intermediate value $\frac{1}{2}$, a process should set a during a 0 to 1 transition and clean it after setting I successfully. Any other process will also set I if it sees that a is set. Once the indicator is set, it can safely clear a to prevent unnecessary future writes to the indicator.

```

shared variables:
 $X = (c, v) : (\mathbb{N} \cup \{\frac{1}{2}\}, \mathbb{N})$ ; initially(0, 0)
parent: scalable indicator
Arrive
  succ ← false
  undoArr ← 0
  while( $\neg$ succ)
     $x \leftarrow$  Read( $X$ )
    if  $x.c \geq 1$  then
      if CAS( $X, x, (x.c + 1, x.v)$ ) then
        succ ← true
    if  $x.c = 0$  then
      if CAS( $X, x, (\frac{1}{2}, x.v + 1)$ ) then
        succ ← true
         $x \leftarrow (\frac{1}{2}, x.v + 1)$ 
    if  $x.c = \frac{1}{2}$  then
      parent.Arrive
      if  $\neg$ CAS( $X, x, (1, x.v)$ ) then
        undoArr = undoArr + 1
  while(undoArr > 0) do
    parent.Depart
    undoArr = undoArr - 1
Depart
  while(true) do
     $x \leftarrow$  Read( $X$ )
    if CAS( $X, x, (x.c - 1, x.v)$ ) then
      if  $x.c = 1$  then parent.Depart
    return

```

Figure 2. Code for hierarchical SNZI node

3 Linearizability

Linearizability [7] is a safety property of concurrent systems. It is formalized as follows.

In a shared memory model \mathcal{M} , $O = \{o_1, \dots, o_k\}$ denotes the set of k shared objects, $P = \{p_1, \dots, p_n\}$ denotes the set of n processes accessing the objects. Shared objects support a set of *operations*, which are pairs of invocations and matching responses. Every shared object has a set of states that it could be in. A *sequential specification* of a (deterministic) shared object is a function that maps every pair of invocation and object state to a pair of response and a new object state.

The behavior of \mathcal{M} is defined as H , the set of all possible sequences of invocations and responses together with the initial states of the objects. A history $\sigma \in H$ induces an irreflexive partial order $<_\sigma$ on operations such that $op_1 <_\sigma op_2$ if the response of op_1 occurs in σ before the invocation of op_2 . Operations in σ that are not related by $<_\sigma$ are concurrent. σ is sequential iff $<_\sigma$ is a strict total order. Let $\sigma|_i$ be the projection of σ on process p_i , which is the subsequence of σ consisting of all invocations and re-

```

shared variables:
 $X = (c, a, v) : (\mathbb{N}, \text{boolean}, \mathbb{N})$ ; initially(0, false, 0)
 $I : \text{boolean}$ ; initially false
Arrive
  repeat
     $x \leftarrow \text{Read}(X)$ 
    if  $x.c = 0$  then  $x' \leftarrow (1, \text{true}, x.v + 1)$ 
    else  $x' \leftarrow (x.c + 1, x.a, x.v)$ 
  until  $\text{CAS}(X, x, x')$ 
  if  $x'.a$  then
     $\text{Write}(I, \text{true})$ 
     $\text{CAS}(X, x', (x'.c, \text{false}, x'.v))$ 
Depart
  repeat
1.    $x \leftarrow \text{Read}(X)$ 
2.   if  $\text{CAS}(X, x, (x.c - 1, \text{false}, x.v))$  then
3.     if  $x.c \geq 2$  then
4.       repeat
5.          $\text{LL}(I)$ 
6.         if  $\text{Read}(X).v \neq x.v$  then return
7.         if  $\text{SC}(I, \text{false})$  then return
Query
  return  $\text{Read}(I)$ 

```

Figure 3. Code for SNZI root node

sponses that are performed by p_i . Let $\sigma|_{o_i}$ be the projection of σ on object o_i , which consists of all invocations and responses of operations that are performed on object o_i .

A sequential history σ is *legal* if it respects the semantics of the objects as expressed in their sequential specifications. More specifically, for each object o_i , if s_j is the state of o_i before the j -th operation op_j in $\sigma|_{o_i}$, then the invocation and response of op_j and the resulting new state s_{j+1} of o_i follow the sequential specification of o_i . Given a history σ , a *sequential permutation* π of σ is a sequential history in which the set of operations as well as the initial states of the objects are the same as in σ . The formal definition of linearizability is given as follows.

Linearizability There exists a sequential permutation π of σ such that 1) for each object o_i , $\pi|_{o_i}$ is a legal sequential history (i.e. π respects the sequential specification of the objects), and 2) if $op_1 <_{\sigma} op_2$, then $op_1 <_{\pi} op_2$ (i.e., π respects the real-time ordering of operations).

In every history σ , if we assign increasing time values to all invocations and responses, then every operation can be shrunk to a single time point between its invocation and response such that the operation appears to be completed instantaneously at this time point [3]. This time point for each operation is called its *linearization point*. Linearizability is defined in terms of the invocations and responses

of high-level operations, which are implemented by algorithms on concrete shared data structures in real programs. Therefore, the execution of high-level operations may have complicated interleaving of low-level actions. Linearizability of a concrete concurrent algorithm requires that, despite of complicated low-level interleaving, the history of high-level interface events still has a sequential permutation that respects both the real-time ordering among operations and the sequential specification of the objects. This idea is formally presented in Section 4 using refinement relations.

4 Verification via Refinement Checking

We model concurrent systems using a process algebra, whose behavior is described using a labeled transition system. Linearizability is then defined as a refinement relation from an implementation model to a specification model.

4.1 Modeling Language

We introduce the relevant subset of syntax of CSP (Communicating Sequential Processes) [8] extended with shared variables. We choose this language because of its rich set of operators for concurrent communications.

Process A process P is defined using the grammar:

$$\begin{aligned}
 P ::= & \text{Stop} \mid \text{Skip} \mid e\{\text{program}\} \rightarrow P \mid P \setminus X \mid P_1; P_2 \\
 & \mid P_1 \square P_2 \mid \text{if}(b) \{P_1\} \text{ else } \{P_2\} \mid P_1 \parallel P_2
 \end{aligned}$$

where P, P_1, P_2 are processes, e is a name representing an event with an optional sequential program program , X is a set of events, and b is a Boolean expression.

Stop is the process that communicates nothing, also called deadlock. $\text{Skip} = \checkmark \rightarrow \text{Stop}$, where \checkmark is the termination event. Event prefixing $e \rightarrow P$ performs e and afterwards behaves as process P . If e is attached with a sequential program, the valuation of the shared variables is updated accordingly. For simplicity, assignments are restricted to update only shared variables. Process $P \setminus X$ hides all occurrences of events in X . An event is invisible iff it is explicitly hidden by the hiding operator $P \setminus X$. Sequential composition, $P_1; P_2$, behaves as P_1 until its termination and then behaves as P_2 . External choice $P_1 \square P_2$ is solved only by the occurrence of a visible event. Conditional choice $\text{if}(b) \{P_1\} \text{ else } \{P_2\}$ behaves as P_1 if the Boolean expression b evaluates to true, and behaves as P_2 otherwise. Indexed interleaving $P_1 \parallel P_2$ runs all processes independently except for communication through shared variables. The generalized form of interleaving is written as $\parallel x : \{0..n\} @ P(x)$. Processes may be recursively defined, and may have parameters (see examples later). The formal operational semantics of our language is presented in [9].

To model nonblocking algorithms, our language provides strong support for synchronization primitives, such as *compare-and-swap (CAS)* and *load-linked (LL)/store-conditional (SC)*, which are elaborated as follows.

CAS¹ The operational semantics of conditional choice require that the condition evaluation and the first event to be executed of true/false branch be finished in one atomic step. Hence CAS primitive can be directly modeled using conditional choice.

```

/* The pseudo code of CAS semantics */
bool CAS(ref addr, val exp, val new) :
  atomically {
    if(*addr == exp) {*addr := new; }
    else { }
  }
/* The CSP representation of CAS */
if(*addr == old)
  {τ{*addr == new; } → Skip}
else {Skip}

```

LL/SC² A shared counter *counter* is added to indicate the timestamp when the content of a memory location X is modified and a counter flag is associated with each of the processes. When LL is executed by one of the processes, the content of X is read and the value of *counter* is stored in the counter flag. If an external event updates X or the process executes an operation that may invalidate an atomic sequence (e.g., an exception), then *counter* is increased by 1. When the corresponding SC is executed, the counter flag is checked. If the flag is equal to *counter*, then SC will be successfully executed. Otherwise, nothing can be done.

```

/* flag[i] denotes the counter flag of process i */
LL(i)    = τ{READ X; flags[i] = counter; } → Skip;
SC(i, v) = if(flags[i] == counter)
          {τ{X = v; counter++; } → Skip}
          else Skip;
Update(v) = τ{UPDATE X; counter++; } → Skip;

```

The semantics of a model is defined using labeled transition systems (LTS). Let Σ denote the set of all visible events and τ denote the set of all invisible events. Let Σ^* be the set of finite traces. Let Σ_τ be $\Sigma \cup \tau$. A LTS is a 3-tuple $L = (S, init, T)$ where S is a set of states, $init \in S$ is the initial state, and $T \subseteq S \times \Sigma_\tau \times S$ is a labeled transition relation. Let s, s' be states in S and $e \in \Sigma_\tau$, we write $s \xrightarrow{e} s'$ to denote $(s, e, s') \in T$. We write $s \xrightarrow{e_1, e_2, \dots, e_n} s'$ iff there exists $s_1, \dots, s_{n+1} \in S$ such that $s_i \xrightarrow{e_i} s_{i+1}$ for

¹CAS atomically compares the content of a memory location to an expected value, and if they are the same, the content of that memory location is assigned to the new given value.

²LL/SC are a pair of instructions. LL first reads the current content from a memory location X . A subsequent SC stores a new value to X only if no updates have happened in between LL and SC; otherwise, it fails.

all $1 \leq i \leq n$, $s_1 = s$ and $s_{n+1} = s'$. Let $tr : \Sigma^*$ be a sequence of visible events. $s \xrightarrow{tr} s'$ iff there exists $e_1, e_2, \dots, e_n \in \Sigma_\tau$ such that $s \xrightarrow{e_1, e_2, \dots, e_n} s'$. The set of traces of L is $traces(L) = \{tr : \Sigma^* \mid \exists s' \in S, init \xrightarrow{tr} s'\}$. In this paper, we consider only LTSs with a finite number of states. In particular, we bound the sizes of variable domains by constants.

Theorem 1 (Refinement). Let $L_{im} = (S_{im}, init_{im}, T_{im})$ be a LTS for an implementation. Let $L_{sp} = (S_{sp}, init_{sp}, T_{sp})$ be a LTS for a specification. L_{im} refines L_{sp} , written as $L_{im} \sqsubseteq_T L_{sp}$, iff $traces(L_{im}) \subseteq traces(L_{sp})$.

4.2 Linearizability

This section briefly shows how to create high-level linearizable specifications and how to use refinement relation to define linearizability of concurrent implementations.

We define the linearizable specification LTS $L_{sp} = (S_{sp}, init_{sp}, T_{sp})$ for a shared object o in the following way. Every execution of an operation of o on a process includes three atomic steps: the invocation action, the linearization action, and the response action. The linearization action performs the computation based on the sequential specification of the object. All the invocation and response actions are visible events, while the linearization ones are invisible events. Their complete specification and transition rules in LTS is presented in [9]. We now consider a LTS $L_{im} = (S_{im}, init_{im}, T_{im})$ that supposedly implements object o . Theorem 2 characterizes linearizability of the implementation through refinement relations.

Theorem 2. Traces of L_{im} are linearizable iff $L_{im} \sqsubseteq_T L_{sp}$.

The proof of theorem 2 is given in [9]. The theorem shows that to verify linearizability of an implementation, it is necessary and sufficient to show that the implementation LTS is a refinement of the specification LTS as we defined above. This provides the theoretical foundation of our verification of linearizability. Notice that the verification by refinement given above does not require identifying low-level actions in the implementation as linearization points, which is the difficult (and sometimes even impossible) task. In fact, the verification can be automatically carried out without any special knowledge about the implementation beyond knowing the implementation code.

5 SNZI Model

In order to prove that SNZI algorithm is a linearizable implementation, we model its specification and implementation in CSP, and then verify whether the implementation refines the specification.

```

ArriveA(i) = arrive_inv.i → τ{surplus++; }
            → arrive_res.i → Skip;
DepartA(i) = depart_inv.i → τ{surplus--; }
            → depart_res.i → Skip;
QueryA()   = query.(surplus > 0) → QueryA();
ProcessA(i) = ArriveA(i); DepartA(i); ProcessA(i)
SNZIA()    = (||| x : {0..P-1}@ProcessA(x))\{τ}
            ||| QueryA();

```

Figure 4. abstract specification model

```

ArriveI(p, n) = arrive_inv.p →
                if(n == 0) ArriveR(p) else Arrive(p, n);
                arrive_res.p → Skip;
DepartI(p, n) = depart_inv.p →
                if(n == 0) DepartR(p) else Depart(p, n);
                depart_res.p → Skip;
Process(i)    = □ x : {0..N-1}@
                (ArriveI(i, x); DepartI(i, x));
Query()      = query.I → Query();
SNZI()       = (||| x : {0..P-1}@Process(x))\{τ}
                ||| Query();

```

Figure 5. concrete implementation model

Fig. 4 shows the abstract specification model with P processes. Each *Arrive* and *Depart* process consists of invocation event, linearization event τ and response event. *Query* process recursively reads whether *surplus* is greater than zero or not. *ProcessA* models the behavior of a process visiting repeatedly where a visit is an *Arrive* followed by a *Depart*. *SNZIA* interleaves every *ProcessA* and hides the linearization events.

The basic structure of the implementation (skip the details of *Arrive* and *Depart* operations) is showed in Fig. 5. To initialize the rooted tree in the implementation, we create a width N array named *node* storing SNZI objects. The root is *node*[0], and for $0 < i < N$, the parent of *node*[i] is *node*[$\lfloor \frac{i-1}{2} \rfloor$]. Since P processes may visit the same node concurrently, an $N \times P$ array is introduced to store the local variables within an operation of P processes visiting N nodes. The full implementation model is in Appendix.

A process could visit any node at any time, that is, which node a process chooses to visit is decided by external environment. Thus, external choice \square is used to represent a process visiting a node randomly. *ArriveI* represents the process p arriving at the node n . If $n = 0$ (the visiting node is the root), then it starts process *ArriveR* which describes a process enters the root. Otherwise, it starts process *Arrive* which is used for a process arriving a hierarchical node. So does *DepartI*. Due to space constraints, we show the resulting code only for *Depart* operation at the

```

1. DepartR(p) = τ{cc[p] = c[0]; aa[p] = a; vv[p] = v[0]; } →
2. if(cc[p] == c[0] && aa[p] == a3. && vv[p] == v[0])
4. {τ{c[0] = cc[p] - 1; a = false; v[0] = vv[p]; } →
5. if(cc[p] > 1){τ → Skip}
6. else{τ → DepartLoop(p)}
7. }else{τ → DepartR(p)};

8. DepartLoop(p) = τ{counts[p] = count; } →
9. if(vv[p] != v[0]) {τ → Skip}
10. else{if(counts[p] != count){τ → DepartLoop(p)}
11. else{τ{I = false; count++; } → Skip}
12. };

```

Figure 6. Depart operation on root node

root in Fig. 6. Since the original algorithm of *Depart* includes two-fold loop statements, each loop is organized as a recursively defined process. *DepartR* process is the entry and the outer loop of the operation, while *DepartLoop* process denotes the inner loop. The original X and x are both structured variables composed of three simple variables, so (C, A, V) and (c, a, v) are used to represent them respectively. An atomic and invisible event τ containing the assignment statements of c, a and v represents the assignment of x on line 1. Similar is X on line 4. For line 5, 6 and 7, another τ is added between if/else condition and the first event of true/false branch to avoid them executed in one atomic step. *DepartLoop* contains a pair of *LL/SC* primitives. The value of *counter* is recorded when performing *LL* (line 8). Then when the process attempts *SC*, it checks whether the recorded value is equal to the current value of *counter* (line 10). If they are not equal, *DepartLoop* is repeatedly invoked (line 10). Otherwise, the process assigns *false* to I and then performs *Skip* event to return the invoking process (line 11).

6 Verification and Experimental Result

Based on Theorem 2, automatic refinement checking allows us to verify the linearizability of SNZI algorithm. PAT [11] supports different notions of refinements based on different semantics. A refinement checking algorithm (inspired by the one implemented in FDR [10] but extended with partial order reduction) is used to perform refinement checking on-the-fly. The key idea is to establish a (weak) simulation relationship from the specification to the implementation. We remark that FDR does not support shared variables/arrays, and therefore, is not easily applicable. Another candidate tool is the SPIN model checker, which supports verification of LTL properties. Nonetheless, formalization linearizability as LTL formulae results in large LTL formulae and thus not feasible for verification.

We have experimented SNZI on PAT for different number of processes and tree nodes. The table below summarizes the results, where '-' means infeasible, and 'POR' means partial order reduction. The testbed is a PC with 2.83GHz Intel Q9550 CPU and 4 GB memory.

Setting		Result without POR		Result with POR	
#Proc	#Node	Time(sec)	#States	Time(sec)	#States
2	2	23.3	28163	17.1	23828
2	3	73.6	62753	41.4	52779
2	4	393	376342	157	173694
2	5	1298	712857	322	341845
2	6	-	-	496	485156
3	2	-	-	6214	8451568

The number of states and running time increase rapidly with data size, and especially the number of processes. This conform to theoretical results [1]: model checking linearizability is in EXPSPACE. We have employed several optimization techniques to improve scalability. First, we use partial order reduction and effectively reduce the search space and running time. Second, we manually combined sequences of process-local statements into atomic blocks, such as organizing consecutive events which only cope with local variables into one single event τ . Third, we specified every operation using a minimum number of processes, in order not to generate multiple equivalent states as different parameterized processes containing the same events.

7 Related Work and Conclusion

The idea of refinement has been explored by Alur, et al. [1] to show that linearizability can be cast as containment of two regular languages. Our definition of linearizability on refinement is more general, regardless of the modeling language and knowledge of linearization points.

Formal verification of linearizability is a much studied research area. There are various approaches in the literature. Verification using theorem provers is another approach [5], where algorithms are proved to be linearizable by using simulation between input/output automata modeling the behavior of an abstract set and the implementation. However, theorem prover based approach is not automatic. Conversion to IO automata and use of PVS require strong expertise. Wang and Stoller [14] present a static analysis that verifies linearizability for an unbounded number of threads. Their approach detects certain coding patterns, hence is not complete (i.e., not applicable to SNZI algorithm). Amit et al. [2] presented a shape difference abstraction that tracks the difference between two heaps. The main limitation of this approach is that users need to provide linearization points, which is generally unknown. A buggy design may have no linearization points at all. In [13], Vechev and Yahav provided two methods for linearizability checking. One method requires algorithm-specific user annotations for linearization points. The other is fully automatic

but inefficient (The worse case time is exponential in the length of the history). As a result, the number of operations they can check is only 2 or 3. In contrast, our approach handles all possible interleaving of operations given sizes of the shared objects.

In this work, we expressed linearizability using refinement relation. By using this definition, we have successfully verified the SNZI algorithms for the first time. We have shown that the refinement checking algorithm behind PAT allows us to successfully verify complicated concurrent algorithms without the knowledge of linearization points. During the analysis, we have faced the infamous state explosion problem. In future, we will explore how to combine different state space reduction techniques and parameterized refinement checking for infinite number of processes.

References

- [1] R. Alur, K. Mcmillan, and D. Peled. Model-checking of correctness conditions for concurrent objects. In *LICS 96*, pages 219–228. IEEE, 1996.
- [2] D. Amit, N. Rinetzky, T. Reps, M. Sagiv, and E. Yahav. Comparison under abstraction for verifying linearizability. In *CAV 07*, pages 477–490. Springer, 2007.
- [3] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics*. John Wiley & Sons, Inc., Publication, 2nd edition, 2004.
- [4] J. Derrick, G. Schellhorn, and H. Wehrheim. Proving linearizability via non-atomic refinement. In *IFM 07*, pages 195–214, 2007.
- [5] S. Doherty, L. Groves, V. Luchangco, and M. Moir. Formal verification of a practical lock-free queue algorithm. In *FORTE 04*, pages 97–114. Springer, 2004.
- [6] F. Ellen, Y. Lev, V. Luchangco, and M. Moir. Snzi: scalable nonzero indicators. In *PODC 07*, pages 13–22, 2007.
- [7] M. Herlihy and J. M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, 1990.
- [8] C. A. R. Hoare. *Communicating Sequential Processes*. International Series in Computer Science. Prentice-Hall, 1985.
- [9] Y. Liu, W. Chen, Y. A. Liu, and J. Sun. Model Checking Linearizability via Refinement. Technical Report TR08c, National Univ. of Singapore, Dec 2008. <http://www.comp.nus.edu.sg/~pat/report1.pdf>.
- [10] A. W. Roscoe. *The Theory and Practice of Concurrency*. Prentice-Hall, 1997.
- [11] J. Sun, Y. Liu, and J. S. Dong. Model checking csp revisited: Introducing a process analysis toolkit. In *ISoLA 2008*, pages 307–322. Springer, 2008.
- [12] J. Sun, Y. Liu, J. S. Dong, and H. H. Wang. Specifying and verifying event-based fairness enhanced systems. In *ICFEM 08*, Oct 2008.
- [13] M. Vechev and E. Yahav. Deriving linearizable fine-grained concurrent objects. In *PLDI 08*, pages 125–135, 2008.
- [14] L. Wang and S. Stoller. Static analysis of atomicity for programs with non-blocking synchronization. In *PPoPP 05*, pages 61–71. ACM Press New York, NY, USA, 2005.

Appendix: The Complete Model of SNZI Algorithms

```

#define P 3; //number of processes
#define N 3; //number of nodes
//-----Shared Variables-----
var c[N]; //X.c in the original algorithm
var v[N]; //X.v
var a = false; //X.a for root node
var I = false; //presence indicator
//-----Local Variables-----
var cc[N * P]; //x.c
var vv[N * P]; //x.v
var aa[P]; //x.a
var rc[P]; //x'.c
var ra[P]; //x'.a
var rv[P]; //x'.v
var s[N * P]; //succ
var u[N * P]; //undoArr
//for LL/SC Primitives
var count;
var counts[P];

//*****The Concrete Implementation Model*****

ArriveI(p, n) = arrive_inv.p → ArriveGeneral(p, n); arrive_res.p → Skip;
DepartI(p, n) = depart_inv.p → DepartGeneral(p, n); depart_res.p → Skip;
ArriveGeneral(p, n) = if(n == 0) ArriveR(p) else Arrive(p, n);
DepartGeneral(p, n) = if(n == 0) DepartR(p) else Depart(p, n);

//===== Start : This part is for root node =====

//-----Arrival at Root Node-----
ArriveR(p) = τ{cc[p] = c[0]; aa[p] = a; vv[p] = v[0]; } → //x ← Read(X)
            if(cc[p] == 0){τ → τ{rc[p] = 1; ra[p] = true; rv[p] = vv[p] + 1; } →
                if(cc[p] == c[0]&&aa[p] == a&&vv[p] == v[0]){τ{c[0] = rc[p]; a = ra[p]; v[0] = rv[p]; } → R2(p)}
                else{ArriveR(p)}
            }else{τ → τ{rc[p] = cc[p] + 1; ra[p] = aa[p]; rv[p] = vv[p]; } →
                if(cc[p] == c[0]&&aa[p] == a&&vv[p] == v[0]){τ{c[0] = rc[p]; a = ra[p]; v[0] = rv[p]; } → R2(p)}
                else{ArriveR(p)}
            };

R2(p) = if(ra[p] == true)
        {τ → τ{I = true; count = count + 1; } →
            if(c[0] == rc[p]&&a == ra[p]&&v[0] == rv[p]){τ{c[0] = rc[p]; a = false; v[0] = rv[p]; } → Skip}
            else{τ → Skip}
        }else{τ → Skip};

//-----Departure from Root Node-----
DepartR(p) = τ{cc[p] = c[0]; aa[p] = a; vv[p] = v[0]; } →
            if(cc[p] == c[0]&&aa[p] == a&&vv[p] == v[0]) //if CAS(X, x, (x.c - 1, false, x.v))
            {τ{c[0] = cc[p] - 1; a = false; v[0] = vv[p]; } →
                if(cc[p] > 1){τ → Skip} else{τ → DepartLoop(p)}
            }else{τ → DepartR(p)};

DepartLoop(p) = τ{counts[p] = count; } → //LL(I)
                if(vv[p]! = v[0]){τ → Skip}
                else{if(counts[p]! = count){τ → DepartLoop(p)} else{τ{I = false; count = count + 1; } → Skip}};

```

```

// ===== = ===== Start : Hierarchical SNZI Node =====
// ----- - - - - - Arrival at Hierarchical SNZI Node -----
Arrive(p, n) =  $\tau\{s[n * P + p] = false; \} \rightarrow \tau\{u[n * P + p] = 0; \} \rightarrow ArriveL1(p, n); ArriveL2(p, n); Skip;$ 

//  $\frac{1}{2}$  can not be expressed in CSP, therefore the value of X.c will be increased as two times as the original.
ArriveL1(p, n) =  $if(!s[n * P + p])$ 
   $\{ \tau \rightarrow \tau\{cc[n * P + p] = c[n]; vv[n * P + p] = v[n]; \} \rightarrow$ 
     $if(cc[n * P + p] > 1) // if x.c \geq 1 then$ 
       $\{ \tau \rightarrow$ 
         $if(cc[n * P + p] == c[n] \& \& vv[n * P + p] == v[n]) // if CAS(X, x, (x.c + 1, x.v)) then$ 
           $\{ \tau\{c[n] = cc[n * P + p] + 2; v[n] = vv[n * P + p]; \} \rightarrow$ 
             $\tau\{s[n * P + p] = true; \} \rightarrow Case2(p, n)$ 
           $\} else\{ \tau \rightarrow Case2(p, n) \}$ 
         $\} else\{ \tau \rightarrow Case2(p, n) \}$ 
       $\} else\{ \tau \rightarrow Skip \};$ 
Case2(p, n) =  $if(cc[n * P + p] == 0) // if x.c = 0 then$ 
   $\{ \tau \rightarrow$ 
     $if(cc[n * P + p] == c[n] \& \& vv[n * P + p] == v[n]) // if CAS(X, x, (\frac{1}{2}, x.v + 1)) then$ 
       $\{ \tau\{c[n] = 1; v[n] = vv[n * P + p] + 1; \} \rightarrow \tau\{s[n * P + p] = true; \} \rightarrow$ 
         $\tau\{cc[n * P + p] = 1; vv[n * P + p] = vv[n * P + p] + 1; \} \rightarrow Case3(p, n)$ 
       $\} else\{ \tau \rightarrow Case3(p, n) \}$ 
     $\} else\{ \tau \rightarrow Case3(p, n) \};$ 
Case3(p, n) =  $if(cc[n * P + p] == 1)$ 
   $\{ \tau \rightarrow ArriveGeneral(p, (n - 1)/2);$ 
     $if(cc[n * P + p] == c[n] \& \& vv[n * P + p] == v[n])$ 
       $\{ \tau\{c[n] = 2; v[n] = vv[n * P + p]; \} \rightarrow ArriveL1(p, n)$ 
     $\} else\{ \tau \rightarrow \tau\{u[n * P + p] = u[n * P + p] + 1; \} \rightarrow ArriveL1(p, n)$ 
     $\} else\{ \tau \rightarrow ArriveL1(p, n) \};$ 
ArriveL2(p, n) =  $if(u[n * P + p] > 0) \{ \tau \rightarrow DepartGeneral(p, (n - 1)/2); \tau\{u[n * P + p] = u[n * P + p] - 1; \}$ 
   $\rightarrow ArriveL2(p, n) \} else\{ \tau \rightarrow Skip \};$ 

// ----- - - - - - Departure from Hierarchical SNZI Node -----
Depart(p, n) =  $\tau\{cc[n * P + p] = c[n]; vv[n * P + p] = v[n]; \} \rightarrow$ 
   $if(cc[n * P + p] == c[n] \& \& vv[n * P + p] == v[n]) // if CAS(X, x, (x.c - 1, x.v)) then$ 
     $\{ \tau\{c[n] = cc[n * P + p] - 2; v[n] = vv[n * P + p]; \} \rightarrow$ 
       $if(cc[n * P + p] == 2) \{ \tau \rightarrow DepartGeneral(p, (n - 1)/2); Skip \}$ 
     $\} else\{ \tau \rightarrow Skip \}$ 
   $\} else\{ \tau \rightarrow Depart(p, n) \};$ 

Process(i) =  $(\square x : \{0..N - 1\} @ (ArriveI(i, x); DepartI(i, x)));$ 
Query() =  $query.I \rightarrow Query();$ 
SNZI() =  $(\| \| x : 0..P - 1 @ Process(x) \setminus \{ \tau \} \| \| Query());$ 

// ***** Abstract Specification Model *****
var surplus = 0;
ArriveA(i) =  $arrive\_inv.i \rightarrow \tau\{surplus = surplus + 1; \} \rightarrow arrive\_res.i \rightarrow Skip;$ 
DepartA(i) =  $depart\_inv.i \rightarrow \tau\{surplus = surplus - 1; \} \rightarrow depart\_res.i \rightarrow Skip;$ 
ProcessA(i) =  $ArriveA(i); DepartA(i); ProcessA(i);$ 
QueryA() =  $query.(surplus > 0) \rightarrow QueryA();$ 
SNZIAbs() =  $(\| \| x : 0..P - 1 @ ProcessA(x) \setminus \{ \tau \} \| \| QueryA());$ 

// ***** Refinement Checking *****
#assertSNZI() refines SNZIAbs();

```